



US006473421B1

(12) **United States Patent**
Tappan

(10) Patent No.: **US 6,473,421 B1**
(45) Date of Patent: **Oct. 29, 2002**

(54) **HIERARCHICAL LABEL SWITCHING
ACROSS MULTIPLE OSPF AREAS**

(75) Inventor: **Daniel C. Tappan, Boxborough, MA
(US)**

(73) Assignee: **Cisco Technology, Inc., San Jose, CA
(US)**

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/280,872**

(22) Filed: **Mar. 29, 1999**

(51) Int. Cl.⁷ **H04L 12/56**

(52) U.S. Cl. **370/351; 370/399**

(58) Field of Search **370/238, 238.1,
370/351, 397, 399; 709/230, 238, 241**

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,243,342 A	9/1993	Kattemalalavadi et al. .	341/106
5,353,283 A	10/1994	Tsuchiya	370/60
5,394,402 A	2/1995	Ross	370/94.1
5,426,637 A	6/1995	Derby et al.	370/85.13
5,452,294 A	9/1995	Natarajan	370/54
5,491,692 A	2/1996	Gunner et al.	370/85.13
5,500,860 A	3/1996	Perlman et al.	370/85.13
5,519,704 A	5/1996	Farinacci et al.	370/85.13
5,555,256 A	9/1996	Calamvokis	370/60.1
5,561,669 A	10/1996	Lenney et al.	370/60.1

5,623,492 A	4/1997	Teraslinna	370/397
5,650,993 A	7/1997	Lakshman et al.	370/236
5,651,002 A	7/1997	Van Seters et al.	370/392
5,991,300 A	* 11/1999	Tappan	370/392
6,339,595 B1	* 1/2001	Rekhter et al.	370/392

* cited by examiner

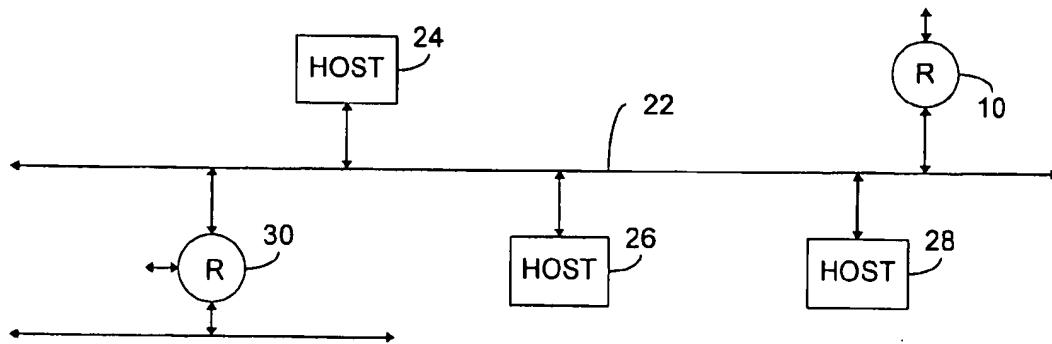
Primary Examiner—Ken Vanderpuye

(74) *Attorney, Agent, or Firm*—Cesari and McKenna, LLP

(57) **ABSTRACT**

In a communications-networking autonomous system consisting of an OSPF domain, autonomous-system border routers (I-ASBR and E-ASBR) cause exchange of hierarchical forwarding labels whose hierarchies are based on OSPF areas. A border router transmits into the domain an OSPF LSA Update message containing an AS-External LSA whose External Route Tag field other routers interpret as specifying a label to be used for forwarding. When that LSA is flooded into the OSPF domain, area border routers respond by flooding new LSAs created from the received one by replacing the label contained in the External Route Tag field with labels that specify their forwarding tables' locations containing information for forwarding to the originating autonomous system border router. In so doing, they enable packets destined for an extra-domain location to be forwarded through the autonomous system without requiring non-border routers to allocate labels to the exterior location or to border routers outside their areas. This enables a networking backbone to conserve its label space even though it carries a large variety of external traffic.

32 Claims, 7 Drawing Sheets



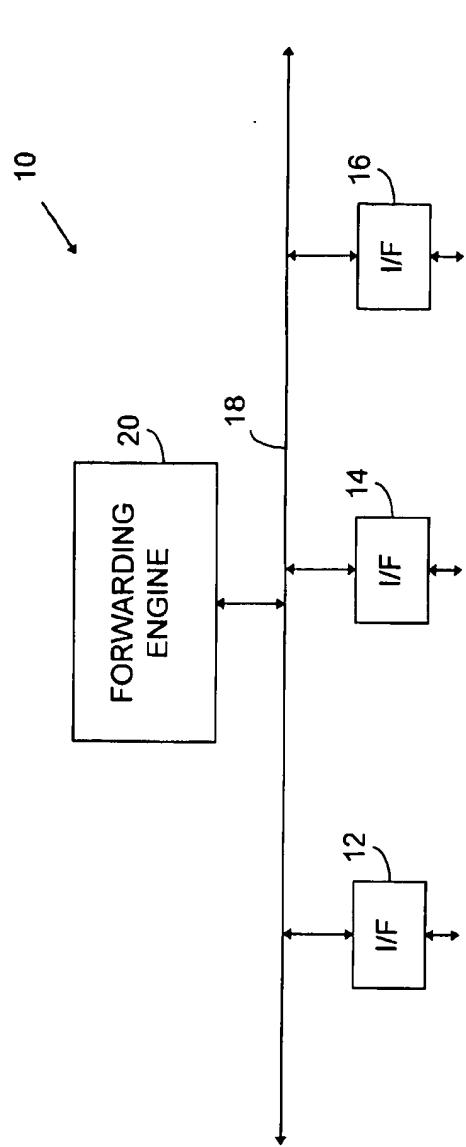


FIG. 1

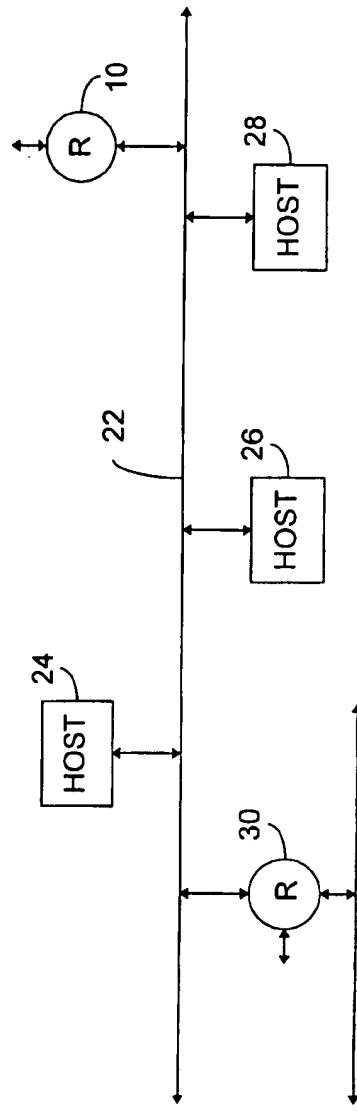


FIG. 2

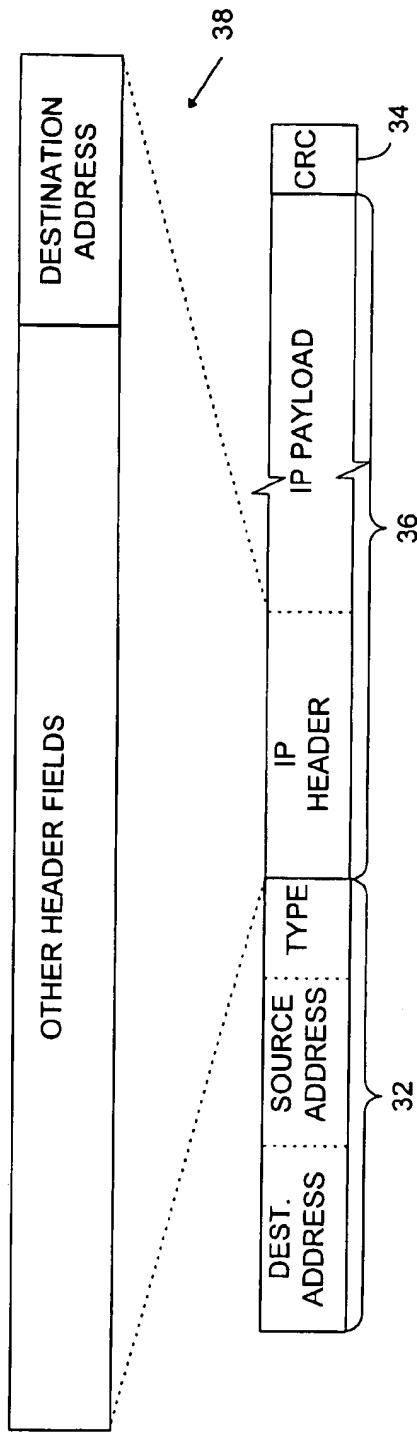


FIG. 3

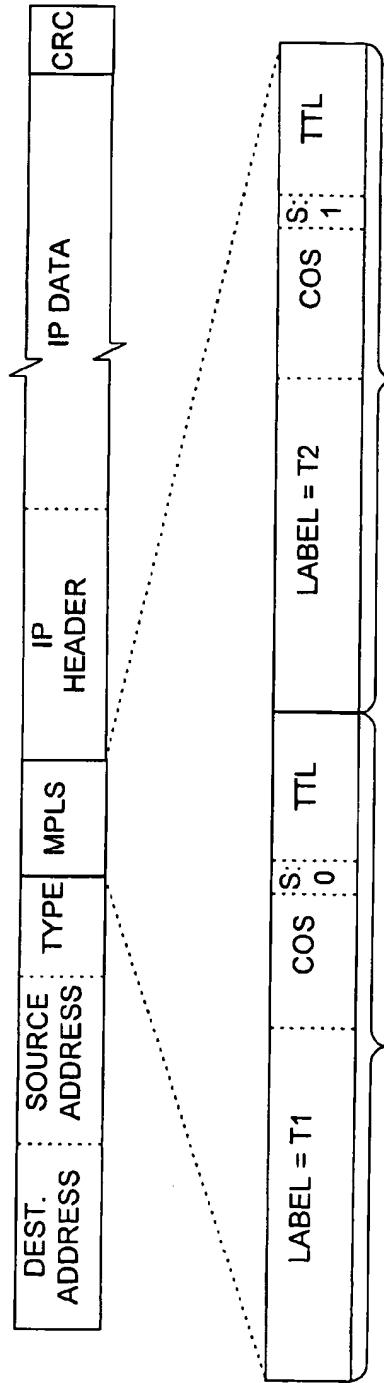
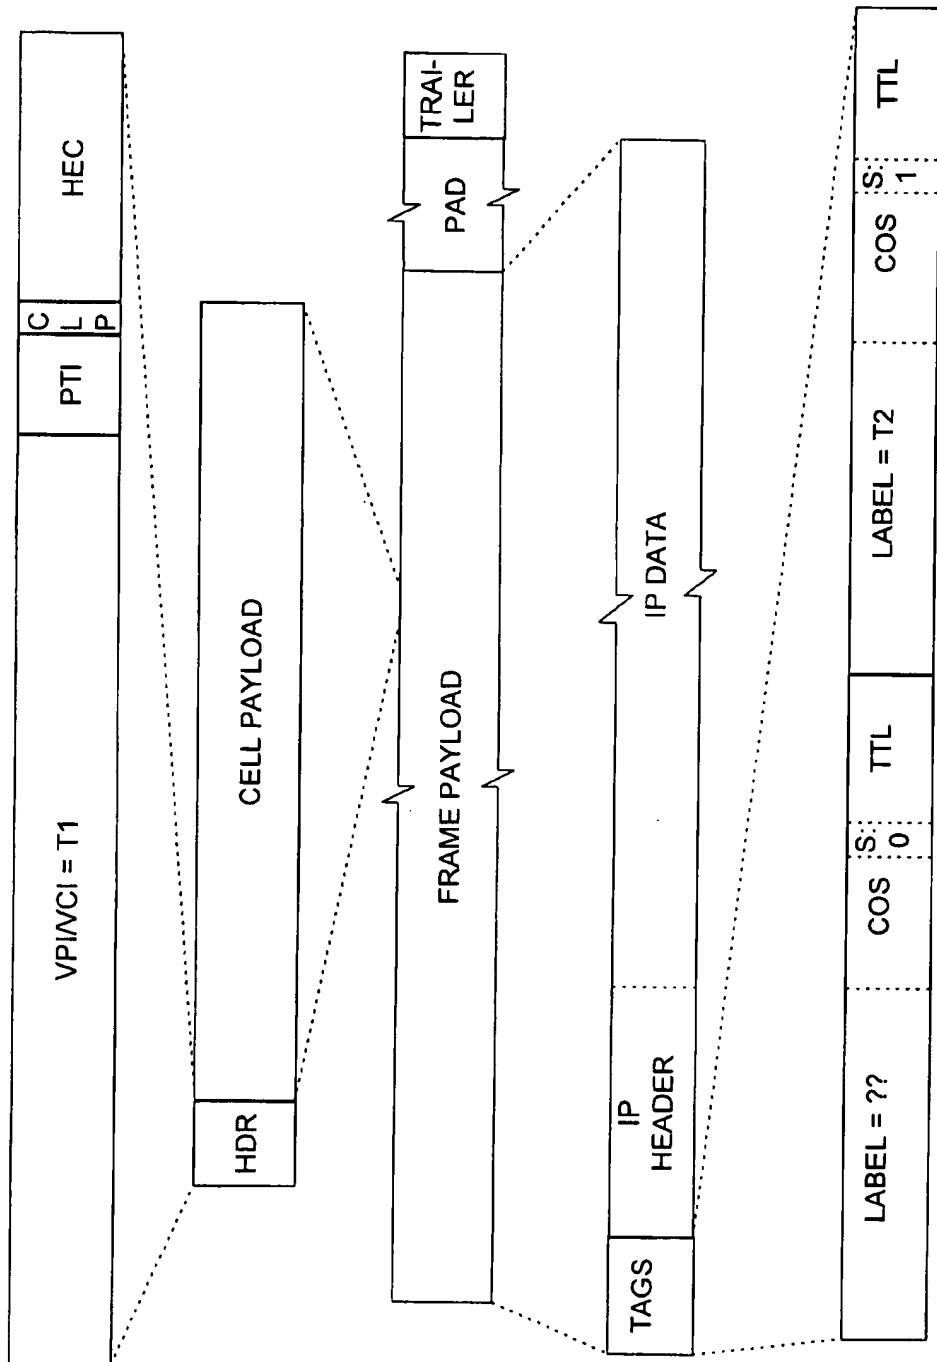


FIG. 4

**FIG. 5**

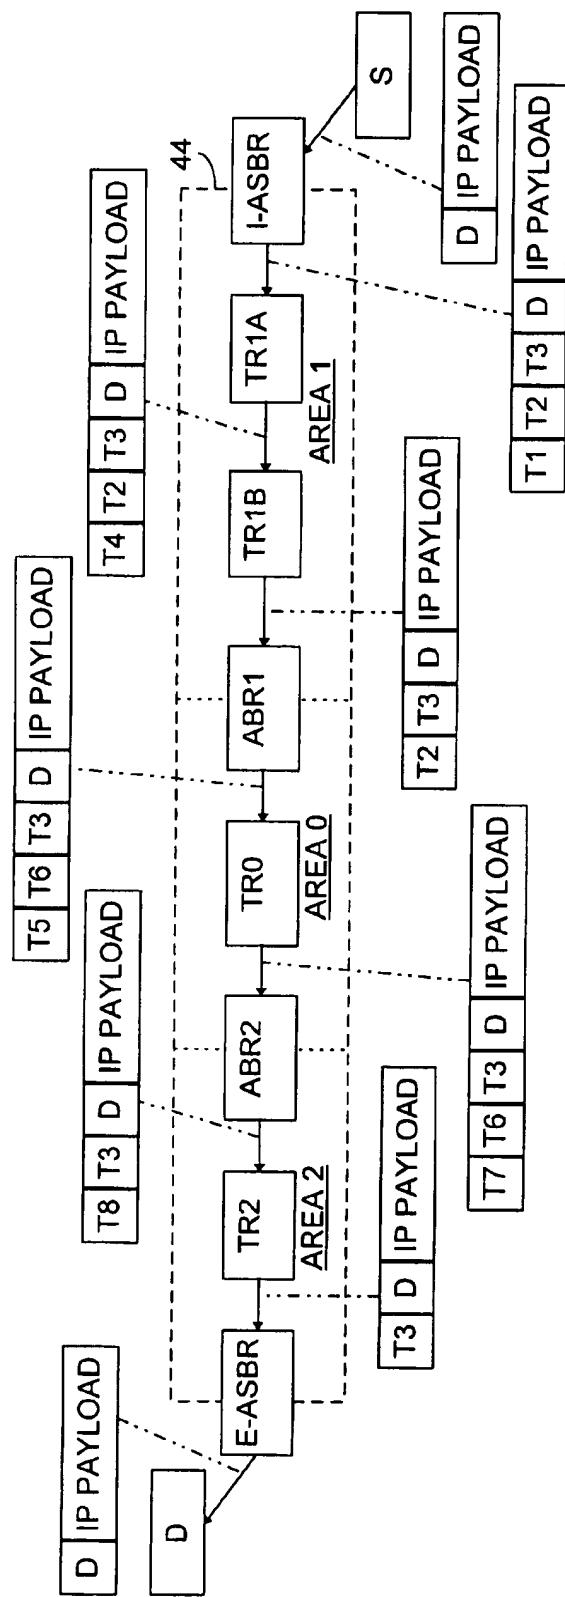


FIG. 6

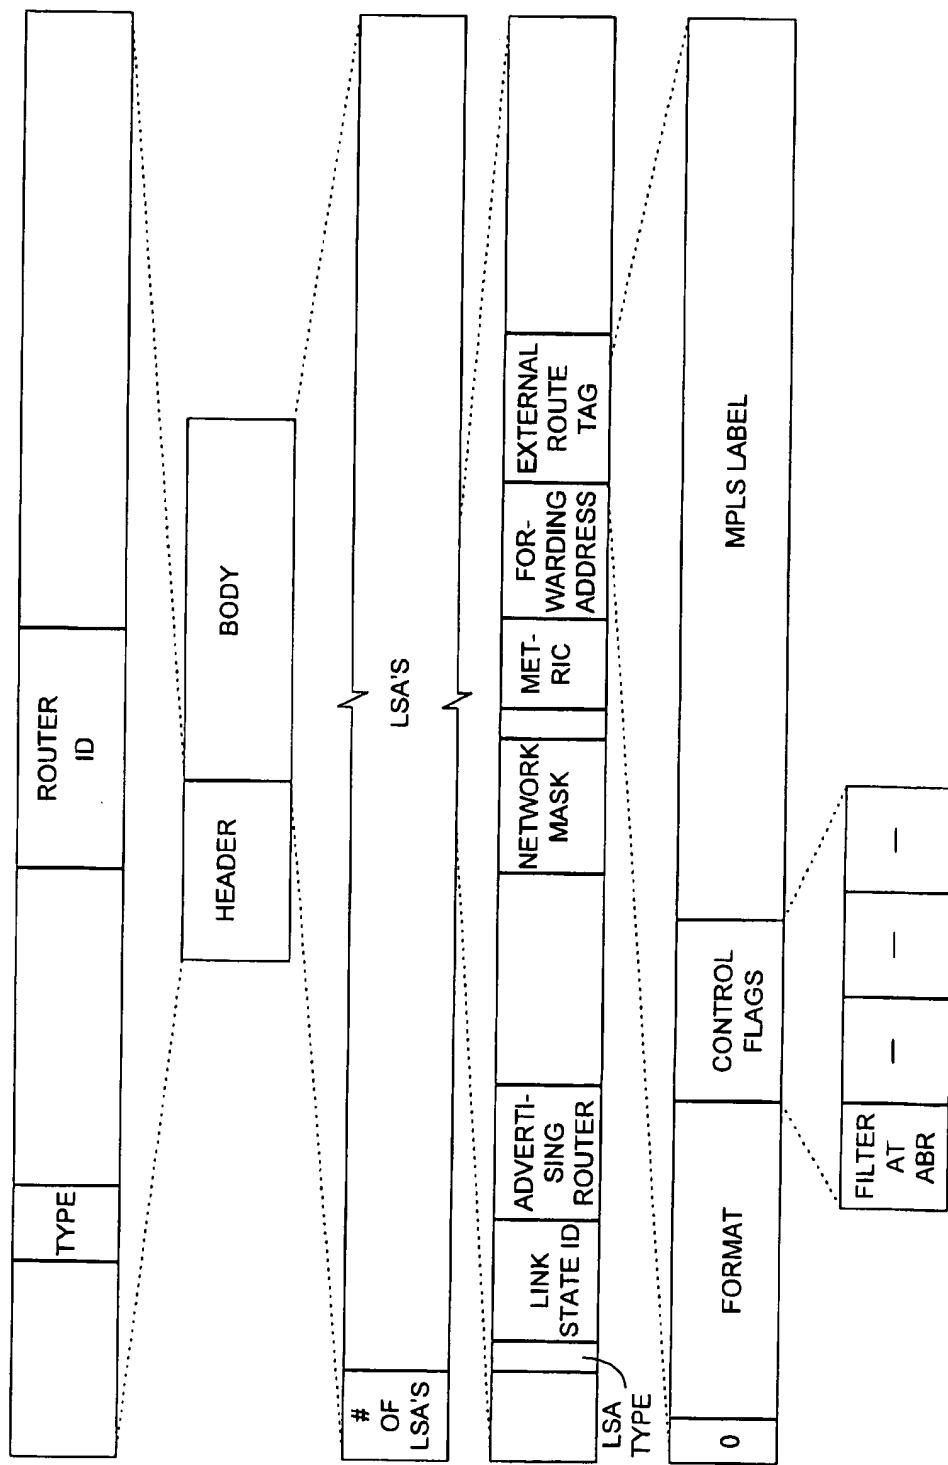


FIG. 7

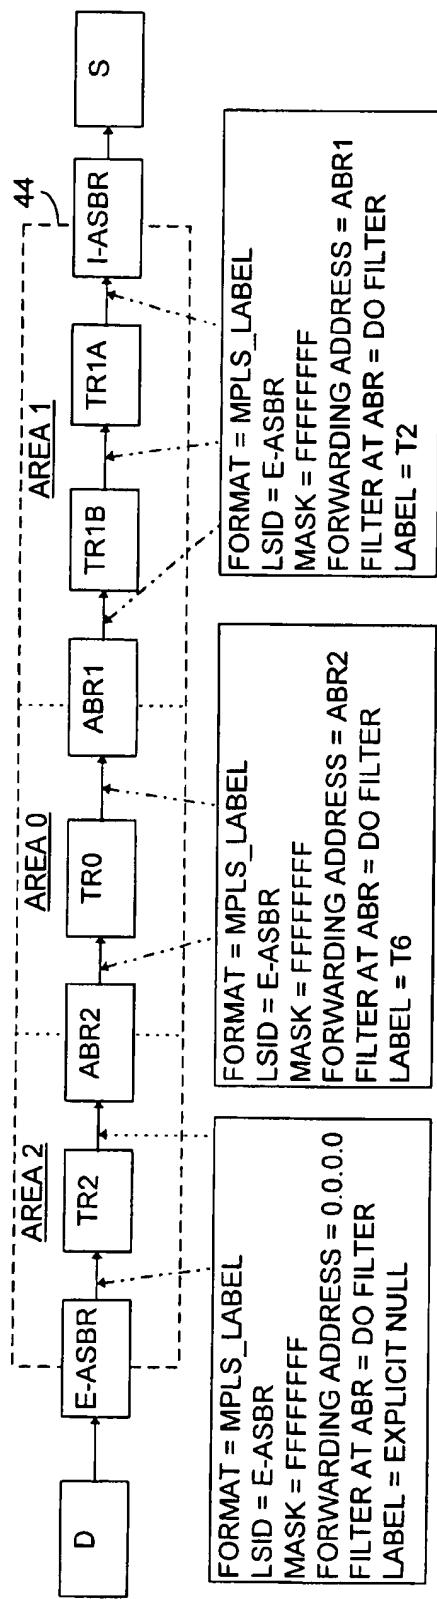


FIG. 8

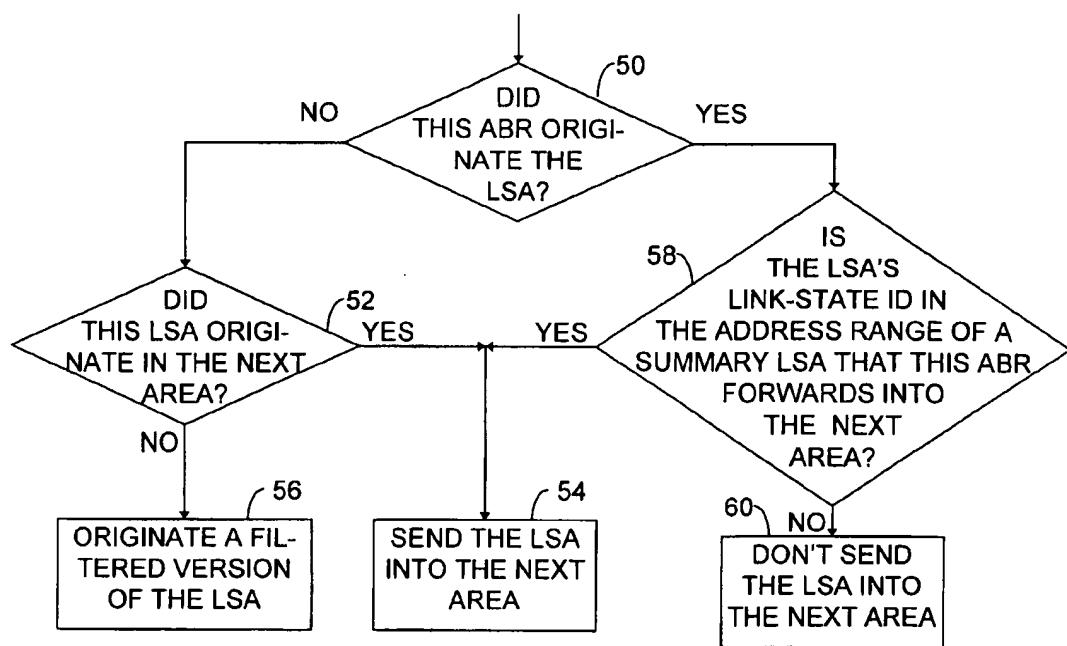


FIG. 9

HIERARCHICAL LABEL SWITCHING ACROSS MULTIPLE OSPF AREAS

BACKGROUND OF THE INVENTION

The present invention is directed to communications networking and in particular to networks that employ label switching.

Internetwork communications are based on operations of routers, which are network devices that determine, on the basis of destination information in packets that they receive, where to forward the packets so that they are likely to reach the intended destinations.

Router configurations vary widely, but FIG. 1 depicts a typical approach. Router 10 includes a plurality of communications interfaces 12, 14, and 16, which send and receive communications packets to and from remote locations. When one of the interface modules receives an incoming packet, it places header information from that packet onto an internal communications bus 18 by which it communicates with a forwarding engine 20, typically a high-performance processor configured by instructions in associated storage circuitry, that determines where the packet should be sent. Once the decision has been made, an output packet is formed from the input packet by packet-assembly circuitry that may reside in one or more of the interface modules and the forwarding engine, and the forwarding engine operates another interface to cause it to send the output packet to a further remote location.

FIG. 2 depicts the router 10 in a local-network environment in which it communicates through one of its interfaces to such remote locations by way of a local-area network bus 22. In a link of that nature, there are typically a number of network devices, such as network devices 24, 26, 28, and 30, that receive the resultant packet-bearing signals, but the packet is not usually intended for all of them. Different systems employ different packet formats to enable their various network devices to distinguish the packets they should read from the ones they should not, but the Ethernet packet format of FIG. 3 is typical.

In that drawing, each packet begins with a link-layer header 32. The link-layer header includes, among other fields, a field that contains a link-layer destination address. If the destination address does not match the address of a network-device interface that receives the packet, that network device ignores the packet.

For present purposes, we will assume that router 10 intends the packet to be received by a further router 30, so the link-layer header's destination-address field will contain the link-layer address of router 30's interface with network link 22. That interface accordingly reads the remainder of the packet, verifying that the contents of a cyclic-redundancy-code trailer 34 are consistent with the remainder of the packet. It then proceeds to process the link-layer packet's payload 36 in accordance with a protocol that the link-layer header's type field specifies.

In the present case, the type field specifies that the link-layer packet's payload is an Internet Protocol ("IP") datagram, which is a network-layer protocol data unit. The purpose of the router's IP process is to determine how to forward the datagram to its ultimate (internetwork-host) destination. To make this determination, the IP process inspects the IP datagram's header 38, and in particular its IP destination-address field. That field's contents identify the host system to which the datagram's contents are to be directed, and router 30 uses this address to determine

through which of its interfaces to forward the packet on to that host system.

The router makes this determination by using a forwarding table, into which it has distilled information about internetwork topology that it has typically obtained by communications with other routers. Routers inform other routers of the host systems to which they can forward communications packets, and they employ such information obtained from other routers to populate their forwarding tables.

Now, the IP address is 32 bits long in most versions and even longer in versions that will soon be adopted, so the IP address could theoretically distinguish among over four billion host systems. Actually, the number of host systems that have globally unique IP addresses is much smaller than this, but the number still is much greater than it is practical for an individual router to have route entries for in its forwarding table.

The solution to this problem historically has been to base the table look-up on destination-address prefixes. That is, some routers will simply indicate that they can handle traffic to all hosts whose destination addresses begin with a particular, say, 16-bit sequence, or "prefix." Additionally, a router may compact its information so as to store routes in this form. Prefixes vary in length, the longest being the most specific and thus presumably representing the best routes to the included host addresses. So when a router receives an IP datagram, it searches through the prefix entries in the forwarding table to find the longest prefix that matches the incoming packet's destination address. When it finds that route in its forwarding table, it reads that route's fields that specify the interface over which it should forward the packet and the link-layer address of the router to which the interface should send the packet for further forwarding.

Although this approach has proved quite serviceable and robust, it has exhibited shortcomings that have led some workers to propose a table-index-based forwarding approach for high-speed networks such as those of some Internet-service providers ("ISPs"). Specifically, routers would inform their neighbor routers of the locations within their tables at which the routes to particular prefixes are located. When their neighbors send them packets destined for those prefixes, those neighbors insert a "shim" between the link-layer header (such as an Ethernet header) and the network-layer header (typically, an IP header). This shim's contents include a label (also referred to as a "tag") that is an index to the desired route in the receiving router's forwarding table.

One of this approach's advantages is that it relieves the receiving router of the need to perform an expensive longest-match search: the label points the receiving router directly to the correct forwarding-table entry. Commonly assigned co-pending U.S. patent application Ser. No. 08/997,343, filed on Dec. 23, 1997 now U.S. Pat. No. 6,339,995, by Rekhter et al. for Peer-Model Support for Virtual Private Networks with Potentially Overlapping Addresses, describes in detail one proposal, known as Multiple-Protocol Label Switching ("MPLS"), for employing such shims. That application also includes an exemplary protocol, referred to as the Tag Distribution Protocol ("TDP") by which routers inform their neighbors of the labels that they associate with their forwarding-table entries. I hereby incorporate that application in its entirety by reference.

FIG. 4 depicts the resultant packet format. The link-layer header has the same format as in FIG. 3, but its type field

identifies the link-layer payload as beginning not with an IP header but rather with an MPLS header, or "shim" between the link-layer header and the network-layer header. FIG. 4 illustrates the MPLS header's format. The MPLS header is organized as a stack of entries, and FIG. 4 gives an example in which there are two entries 40 and 42. In addition to other information, each entry includes a label, which is an index into the forwarding table of the label-switching router that receives it. When a router receives such a packet, it consults the forwarding-table entry that the label specifies and replaces that label with a replacement label that the specified forwarding-table entry contains. That replacement label is typically one that the next router on the path to the requested destination has asked to be included in packets sent to it and intended for the destination with which the forwarding table is associated. For reasons that will become apparent below, the MPLS header may contain more than one label, and the end-of-stack ("S") bit in a stack entry indicates whether it is the bottom entry in the stack. That bit is not set in FIG. 4's entry 40, so that the stack entry is not the bottom one, but it is set in entry 42, which therefore is the bottom stack entry.

Although the formats described in FIGS. 3 and 4 are typical formats for packets exchanged between label-switching routers, they are not the only formats that such routers employ. Routers that communicate with each other over a point-to-point link, i.e., not by way of a shared medium, typically would employ a link-level protocol different from the Ethernet protocol just described, and the formats employed on some "Ethernet" links are actually somewhat more complicated than the format depicted here. In addition, packet forwarding often occurs by way of Asynchronous Transfer Mode ("ATM") links.

In the case of label-switching routers implemented in ATM switches, the IP datagrams are encapsulated in ATM frames. FIG. 5's third row depicts an ATM frame, and its fourth and fifth rows show that the frame's payload is similar to the IP datagram and shim header that FIG. 4's Ethernet header and trailer encapsulate. The only difference is that FIG. 5's fifth row represents the top label by question marks, which indicate that the top label's contents do not matter.

The reason why they do not is that the routing decisions based on those contents when the label-switching router is implemented as a conventional IP router are instead based on an ATM VPI/VCI field in the header of an ATM "cell" when the label-switching router is implemented as an ATM switch. From the point of view of an ATM client, the frame of FIG. 5's third row is the basic unit of transmission, and it can vary in length to as much as 64 Kbytes of payload. (Those skilled in the art will recognize that there are also other possible ATM frame formats, but FIG. 5's third row depicts one, known as "AAL5," that would typically be employed for user data.) From the ATM switch's point of view, though, the basic transmission units are fixed-size cells into which the frames are divided.

Each cell consists of a header and a payload, as FIG. 5's second row illustrates. Among the purposes of the header's PTI field, depicted in FIG. 5's first row, is to indicate whether the cell is the last one in a frame. If it is, its last eight bytes form the frame trailer field that FIG. 5's third row depicts. Among other things, the trailer indicates how much of the preceding cell contents are actual payload, as opposed to padding used to complete a fixed-size cell.

The header field of interest to the present discussion is the VPI/VCI field of FIG. 5's first row. As is well known to those skilled in the art, ATM systems organize their routes into "virtual channels," which may from time to time be

grouped into "virtual paths." Each switch associates a local virtual path/virtual channel indicator (VPI/VCI) with a channel or path that runs through it. When an ATM switch receives a cell, it consults the cell's VPI/VCI field to identify by table lookup the interface through which to forward the cell. It also replaces that field's contents with a value indicated by the table as being the next switch's code for that path or channel, and it sends the resultant cell to the next switch. In other words, the function performed by the VPI/VCI field enables it to serve as the stack's top label. This is why a label-switching router implemented as an ATM switch can ignore the top label field, on which other implementations rely.

It is desirable to use ATM switches in high-volume, "backbone" routers: their fixed cell sizes and use of virtual-channel indicators make it possible for them to operate very rapidly, as is desirable in high-bandwidth backbone networks. Because of the hardware ordinarily employed to implement them, though, ATM switches are subject to one particular drawback, and that is the size of the potential VCI space: the number of VCIs that a given ATM switch can have committed at any one time is relatively small. When a first ATM switch is to forward to a second ATM switch a packet intended for a new destination, the first switch asks the second switch to allocate a VCI for use by that first switch in sending the second switch any packets intended for that destination. Consequently, if the switch is interconnected with many other switches, all of which need its services to forward packets to many different destinations, the switch may be called on to allocate a number of VCIs that exceeds its capacity. Even in systems that do not employ ATM switches to implement label-switching routers, the memory capacity required for a large number of forwarding-information-base entries can be costly.

Proposals have been made to limit forwarding-table size by employing multi-level labeling. This is a type of tunneling, in which the packet being forwarded includes one or more labels representing intermediate destinations, typically in addition to a "bottom" label representing the ultimate destination. To forward such packets, switches along the route to the intermediate destination need not have allocated labels to the ultimate destination. The Rekhter et al. application described above sets forth an example by which, say, a service provider's edge router—i.e., a router with a link to one of the service provider's customers—attaches a top label requesting the route to another edge router as well as a bottom label representing the destination in the customer's network.

SUMMARY OF THE INVENTION

I have recognized that a further reduction in required label space can be achieved by taking advantage of the area concept employed by Open Shortest Path First ("OSPF"), a protocol popularly employed by routers within an autonomous system to exchange the topological information on which they base their routing decisions. Routers in a commonly administered group of networks conventionally employ OSPF to maintain a consistent view of the topology within that group of routers, which we will refer to as a "routing domain." Although the general intention of OSPF is for all of its routers to maintain a common map of the networks within the routing domain, routers in many existing OSPF domains are sometimes configured to divide the domain into "areas" between which there are topological-information differences. Specifically, certain routers have links to more than one area and are configured to consider themselves "area border routers." Such routers export only

a summary of one area's detailed topological information into another area, and this reduces the amount of topological information the various areas must maintain in order to perform their routing functions adequately.

I have recognized that it is possible to use this existing area division as basis for hierarchical labeling. I have further recognized that it is possible to achieve such labeling in a way that is compatible with the normal OSPF protocol and in fact makes use of it. Specifically, I employ the "External Route Tag" field of the OSPF's AS-External LSA to carry labels that the label-switching routers employ, and I have area border routers respond to such messages by filtering them in such a manner as to "tunnel" across OSPF areas.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention description below refers to the accompanying drawings, of which:

FIG. 1, described above, is a block diagram of a router;

FIG. 2, described above, is a block diagram of a local-area network to which the router may be attached;

FIG. 3, described above, is a format diagram of an Ethernet frame containing an IP datagram as a payload;

FIG. 4, described above, is a format diagram of an Ethernet frame in which an MPLS "shim" is disposed between the Ethernet header and the IP header;

FIG. 5, described above, is a format diagram of an ATM frame containing an MPLS header and IP datagram;

FIG. 6 is a diagram illustrating packet forwarding in accordance with the present invention's teachings;

FIG. 7 is a format diagram of an OSPF LSA-Update message employed in accordance with the present invention's teachings;

FIG. 8 is a diagram illustrating LSA propagation in accordance with the present invention's teachings; and

FIG. 9 is a flow chart that illustrates part of the LSA-forwarding-decision process that an area border router performs in accordance with the present invention's teachings.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

To understand how the present invention reduces routers' required topological informational storage, consider the data-packet-forwarding scenario that FIG. 6 illustrates. FIG. 6 depicts transmission of a packet from a source router S to a destination node D. The forwarding path passes through a routing domain 44 to which neither S nor D belongs. That is, routers within routing domain 44 maintain a consistent view of that domain's internal topology by exchanging OSPF messages with each other, but they send no OSPF messages outside the domain. This may be because the domain 44 is the property of, say, an ISP, whereas nodes S and D are owned by that ISP's customers. (As those skilled in the art will recognize, routers' within domain 44 will typically use, say, the Border Gateway Protocol ("BGP") to send router S some type of routing information. But the information it thereby receives about the domain's internal topology depends on that domain's policy and is normally extremely limited.)

When router S needs to send a packet to router D, it generates an IP packet that FIG. 6 represents in abbreviated form as including its IP payload and the destination IP address D. The first of domain 44's routers' to receive the packet will be referred to as "I-ASBR," for Ingress Autonomous System Border Router. In the terminology of exterior

routing protocols, i.e., those such as BGP that are used to exchange information between domains that are not commonly administered, a domain such as domain 44 is referred to as an "autonomous system," and I-ASBR is an autonomous-system border router because it communicates with nodes outside the autonomous system.

As FIG. 6 indicates, I-ASBR receives a normal IP packet from external router S; the drawing represents that packet by its IP payload and the "D" value in its IP header's destination-address field. Receiving such a non-labeled packet, I-ASBR searches its forwarding table for the entry having the longest-prefix match. Let us say that the longest-matching prefix is prefix P. Now, there are a number of ways in which I-ASBR could have obtained information placed in that entry about how to route a packet to P-prefixed destinations. For the sake of example, we will assume that I-ASBR has received this information in a BGP message sent to it by another autonomous-system border router, E-ASBR. As stated above, BGP is an exterior routing protocol: autonomous-system border routers use it to communicate forwarding information to other autonomous-system border routers. Additionally, border routers in a common autonomous system use that protocol to exchange among themselves topological information concerning exterior routes so that they can provide a consistent appearance to exterior routers. One such BGP message would have been a message from E-ASBR to I-ASBR telling I-ASBR that E-ASBR can forward packets to exterior P-prefixed destinations (including D) and that packets sent through E-ASBR should be labeled with label T3 so as to enable E-ASBR to locate its entry for those destinations in a rapid manner. The Rekhter et al. application mentioned above describes a way of using BGP for this purpose.

In creating its forwarding-table entry for destination D's network P, though, I-ASBR additionally entered other labels in the label stack to be placed on packets forwarded in accordance with that entry. As is conventional, it would typically have included as the label stack's top label a label that its immediate downstream neighbor, TR1A, has requested to be used for certain packets sent to it.

In accordance with the present invention, though, that top label specifies TR1A's forwarding-table location of the entry, not for forwarding to the egress autonomous-system border router E-ASBR, but rather for forwarding to router ABR1, which is a border router in I-ASBR's OSPF area, Area 1. Also in accordance with the present invention, an OSPF message to I-ASBR contained the information that it needed to set up its forwarding-table entry in this manner. As will now be explained, that OSPF message was ultimately caused by one that E-ASBR sent to its upstream neighbor TR2.

FIG. 7 depicts the OSPF message format that was used for this purpose. That A drawing's second row depicts the conventional header/body organization of an OSPF message. Such a message is the payload of an IP datagram: the IP datagram's protocol field contains the code that identifies the payload as an OSPF message. The drawing's top row depicts selected header fields. Of these, the type field indicates which of the several standard OSPF-message formats the message assumes. In this case, the type is "LSA Update," whose purpose is to transmit one or more link-state advertisements ("LSAs") from one router to a (typically) neighbor router. An LSA is the unit of topological information that routers exchange under OSPF. The router that is sending the LSA update is specified in the header's router ID field.

When the OSPF message is of the LSA Update type, the message's body assumes the format that FIG. 7's third row

depicts. The body's first field indicates how many LSAs the update message contains, and the remainder of the body contains those LSAs. Different LSAs have different formats, the particular format assumed by an LSA being indicated by its LSA-type field depicted in FIG. 7's fourth row. The particular LSA format of interest here is identified in the Internet community's Request for Comments No. 2178 ("RFC 2178") as an "AS-External LSA." This is the type of LSA that a router ordinarily employs to transmit information regarding external destinations, i.e., destinations outside the OSPF domain. In accordance with the present invention, though, that type of LSA is employed to install hierarchical labels along routes to E-ASBR, as will now be explained.

E-ASBR sends other routers in its area, including router TR2, an LSA of the type just described. The Link State ID field of FIG. 7's fourth row, which is used in an LSA to specify the router or other topological item about which the LSA reports, contains an address of E-ASBR, while the Network Mask field contains all ones to indicate that the Link State ID field contains a complete IP address rather than an IP address prefix.

The LSA's Advertising Router field identifies the router that initiated the LSA, in this case E-ASBR. (As will be seen, routers that receive the message containing the LSA store some of its contents and then relay the information to other routers. In doing so, they preserve the contents of the Advertising Router field but substitute their own IDs in the Router ID field.) An AS-External LSA is conventionally an indication that the advertising router has become aware of an external destination. To indicate that the advertising router can relay data packets to that external destination, it places the address 0.0.0.0 in the Forwarding Address field; otherwise, that field receives the address of another router, through which access to the external location can be had. Although E-ASBR is here using the AS-External LSA for a purpose other than to report an external destination's location, it still places the default route 0.0.0.0 in its Forwarding Address field, for a reason that will manifest itself presently.

From the point of view of E-ASBR, the purpose of this LSA is to tell upstream (in the FIG. 6 sense) neighbor routers that, as will be illustrated below, they should "pop" the label stack of a packet whose top label represents the upstream neighbor's forwarding-table entry for E-ASBR. To signal its intention that the LSA be interpreted as conveying such label information, E-ASBR does two things. First, it places a zero in the first bit of the LSA's External Route Tag field. When that bit is set, existing OSPF speakers interpret the remainder of that field in accordance with RFC 1403, which describes a way to input into the OSPF domain a route that an autonomous-system border router has learned by way of BGP. By keeping this bit zero, E-ASBR is preventing routers from interpreting the field in that manner. Second, it places an MPLS label in the External Route Tag field. The particular format by which this invention's teachings are implemented to accomplish this is not critical, but FIG. 7's fifth and sixth rows depict a format that can be employed for this purpose. As a practical matter, there should also be some bits allocated to specifying that the format is indeed of the MPLS-label-carrying variety; this permits other interpretations also to be given to this field when its first bit is zero. In the illustrated example, the Format field, which may be, say, seven bits in length, serves this purpose. That is, it contains a predetermined label container-indicating value, which indicates that the remainder of the external Route Tag field is of the MPLS-label-carrying type. This means that it includes an MPLS-label field, as FIG. 7's fifth row illus-

trates. That field, which may be, say, twenty bits in length, contains the label whose advertisement is this particular LSA's purpose.

In the case of the LSA generated by E-ASBR, that field contains the value that in MPLS terminology is referred to as an explicit NULL. In MPLS labeling a NULL label can be interpreted to mean either to remove the top label or not to add additional labels when labeled packets are forwarded to the destination indicated. So the effect of that LSA from E-ASBR is that TR2 and ABR2 should not add an additional label based on that LSA.

In a sense this LSA is usually superfluous, because, as will be explained below, Area 2 routers will not ordinarily refer to it in generating their forwarding tables. But placing it in an OSPF LSA Update message triggers the OSPF flooding behavior that we now consider.

When transit router TR2 receives this AS-External LSA, it stores it for possible use in constructing or updating its forwarding table, and, in accordance with normal OSPF flooding practice, it sends a corresponding LSA in an LSA Update message to ABR2. As FIG. 8 indicates, the LSA that it sends to ABR2 is essentially the same as the one it received from E-ASBR.

In accordance with the present invention, the way in which ABR2 deals with this LSA differs from the way in which TR2 does, because ABR2 is configured as an area border router.

When an OSPF-speaking router receives an LSA from a neighbor router, it determines whether it already has an equally or more current version of that LSA. (OSPF provides sequence numbers by which OSPF speakers can determine which version is later.) If not, the router installs that LSA in a topological-information base from which it calculates its forwarding table. If the router is an area border router of the type that the illustrated embodiment employs, it additionally originates a filtered version of that LSA, as will be described presently. Moreover, the ABR also participates in a "flooding" procedure by sending the received or newly originated LSA to other neighbor routers.

When, as is the case in the scenario that FIG. 8 illustrates, the received or newly generated LSA is an AS-External LSA in which the External Route Tag field's first bit is zero and its Format part identifies it as containing an MPLS label, an area border router such as ABR2 that implements the present invention's teachings executes the FIG. 9 routine to determine whether to send the LSA to a given neighbor router as part of the flooding procedure. In the case of the LSA that ABR2 receives from TR2, the result of the decision represented by block 50 is negative: it was E-ASBR, not ABR2, that originated that LSA. So the routine proceeds to that drawing's block 52, which represents determining whether the LSA originated in the area into which ABR2 is deciding whether to send it. Although FIG. 8 shows no ABR2 neighbor in Area 2 other than TR2, from which ABR2 received that LSA, there could be others, and ABR2 would send the received LSA to them, as block 54 indicates.

When ABR2 is determining whether to send the received LSA to TR2 in Area 0, though, the result of block 52's decision is negative. For that situation, the FIG. 9 chart is based on an assumption concerning the Control Flags field that FIG. 7's sixth row depicts. The illustrated embodiment allocates four bits of the External Route Tag field to control flags. This is somewhat arbitrary, since three of them are left for further enhancements. That is, those bits are ignored in the absence of such enhancements. In accordance with convention, their value should be zero in messages from a router in which such enhancements are not implemented.

In the illustrated embodiment, though, the first of the control flags is a Filter at ABR flag, and in the illustrated scenario that flag is set. Routers, such as TR2, that are not area border routers ignore that flag. But when an area border router implementing the illustrated embodiment of the present invention's teachings receives an AS-External LSA in which the first bit and Format part of the External Route Tag field indicate that the LSA is of the MPLS-label-containing variety, it consults that flag.

The above-mentioned assumption on which FIG. 9 is based is that the Filter at ABR flag is set. When it is, an area border router receiving the LSA interprets it as an instruction not to send, as TR2 did, an essentially identical LSA, i.e., an LSA advertising E-ASBR's requested label for E-ASBR-destined packets. Instead, it advertises a label to be used on packets sent to ABR2 for forwarding to E-ASBR. More generally, an ABR should place in the LSA's MPLS Label field a label that routers should use in forwarding packets by way of the router identified by the LSA's Forwarding Address field to the destination that the LSA's Link State ID represents.

So ABR2 generates the LSA by copying the Link State ID, Network Mask, and Format fields and the Filter at ABR flag from the received LSA, setting the Forwarding Address field to ABR2's address, and placing in the MPLS Label field a label value, say, T6, that identifies the location of ABR2's forwarding-table entry that tells where to forward packets destined for E-ASBR. ABR2 must then decide into which areas to flood the thus-generated LSA. (Since it is originating a new LSA, it also puts its own ID in the Advertising Router field.)

For this filtered LSA, ABR2 is indeed the originator, so the result of the block-50 test is positive. ABR2 therefore makes the decision that block 58 represents. In the case of ABR2, the purpose of that decision is to prevent ABR2 from flooding a translated LSA back into Area 2. ABR2 may receive the MPLS LSAs from any of the areas to which it is connected. It will end up basing the replicant on the "closest" one, but there is no guarantee of the order in which the LSAs will be received.

In particular, if some other ABR (not shown) is also connected to both Area 2 and Area 0, then ABR2 may receive from Area 0 an AS-External LSA for E-ASBR originated by that other router, and that LSA may arrive before the one originated by E-ASBR. In that situation, ABR2 will generate a translation. Flooding that translated LSA into Area 2 would waste OSPF control bandwidth.

To appreciate the way in which the illustrated embodiment avoids this result, recall that routers configured as OSPF area border routers send into one area only summaries of the topological information that describes an adjacent area. More specifically, they send LSAs of a type, known as a Summary LSA, that reports to one area the presence within another area of an address range without specifying the other area's particular addresses within that range. The block-58 test determines whether the LSA's link-state ID (which identifies E-ASBR in this case) specifies a router that has an address in a range for which the ABR sends Summary LSAs into the proposed neighbor area. In the case of an E-ASBR-specifying LSA received from Area 0 for which ABR2 is determining whether to flood it into Area 2, the answer would be no, so the flooding and attendant bandwidth waste would not occur.

But the answer is yes in the particular case that FIG. 8 illustrates; E-ASBR, being in Area 2, would have an interface address within a range that ABR2 is configured to report

into Area 0. ABR2 therefore sends the newly originated AS-External LSA into Area 0. When TR0 receives this LSA, it stores it and sends its neighbor ABR1 an LSA that is essentially identical, just as transit router TR2 did; it does not filter the LSA, as ABR2 did, because TR0 is not an area border router.

On the other hand, ABR1 is an area border router, so it generates a new AS external LSA, which requests use of a label that specifies an entry in ABR1's forwarding table. Let us refer to it as T2 the label whose use it thereby requests. As FIG. 8 indicates, routers TR1B and TR1A receive and send essentially identical LSAs. I-ABR receives the one sent by TR1A. I-ASBR may send LSAs to further routers in that domain in response, but the example topology includes no such other neighbors within the OSPF domain. (Router S is not in OSPF domain 44, so it does not receive OSPF messages from that domain.)

We now return to consideration of what I-ASBR does in response to a data packet from the external source router S. As was mentioned above, I-ASBR performs a longest-match search for the address, D, in the incoming IP packet's destination-address field. As was also mentioned above, it thereby finds an entry installed as a result of an internal BGP message from E-ASBR saying that E-ASBR can afford access to prefix P, which is I-ASBR's longest match for destination D. The BGP message also directs it to label such packets with label T3, which identifies the location in E-ASBR's forwarding table that tells E-ASBR how to forward packets to destinations matching prefix P. In accordance with the just-described OSPF LSA-Update message from TR1A, moreover, I-ASBR has been directed to send packets to E-ASBR by way of ABR1 and to apply the label, T2, by which ABR1 identifies its forwarding-table location that tells how to forward packets to E-ASBR. But ABR1 is not an immediate neighbor of I-ASBR. By normal route-calculation operations, I-ASBR has identified TR1A as the next hop for packets that must go through ABR1. And TR1A, which is also a label-switching router, has sent I-ASBR a TDP message requesting that it attach a label T1 to packets that it should forward to ABR1.

So I-ASBR's forwarding-table entry for prefix P, i.e., for the best match for destination D, directs that it apply to B-destined packets a label stack whose top label is T1, whose middle label is T2, and whose bottom label is T3.

In accordance with normal label-switching procedures, TR1A receives the thus-labeled packet, retrieves the information in its forwarding-table entry specified by the top label T1, replaces it with the replacement label, T4, that its forwarding-table entry contains, and forwards the data packet to a further transit router TR1B. As FIG. 6 indicates, router TR1B removes the top label T4 without replacing it. The reason for this is that a TDP message from ABR1 regarding packets to be forwarded to ABR1 gave the implicitNULL value for the label. Since TR1B has removed T4, it "exposes" to ABR1 the label, T2, that specifies ABR1's forwarding-table location that tells it how to forward packets to E-ASBR.

In preparing that forwarding-table entry, ABR1 employed the above-described LSA from TR0. ABR1 used that LSA in generating its forwarding-table entry in a manner the same as that employed by I-ASBR to construct its forwarding-table entry. More mechanically, when an ABR receives an AS-External LSA message in which the Link State ID field specifies an ASBR, and the External Route Tag field's first bit and Format fields identify it as an MPLS-label-carrying LSA, then a non-NULL value in the MPLS-label field tells

it to place into the forwarding-table entry for the destination represented by link-state ID a two-label replacement-label field. The top label is the label that the ABR has been requested to place in packets being sent to the destination that the LSA's Forwarding Address field specifies. The next label is the one that the LSA's MPLS Label field specifies. In this case, ABR1 uses as the top label a value, T5, that TR0 has asked it, in a TDP message, to use in packets sent to TR0 for forwarding to the router, ABR2, identified in the AS-External LSA's Forwarding Address field. The lower label, T6, is the one that ABR2 placed in the AS-External LSA's External Router Tag field as the label for packets sent through that router to the router, E-ASBR, that the LSA's link state ID identified. So ABR1 replaces T2 with T5 and T6.

Now, none of these operations depends on the manner in which the LSRs have been implemented. As was explained above, the top labels may be contained not only in MPLS headers but additionally, if the label-switching routers are implemented in ATM switches, ATM-cell virtual-channel-indicator fields. For the sake of example, though, we will assume that the Area 1 and Area 2 links are all, say, Ethernet and/or SLIP links, whereas the two links illustrated in Area 0 are ATM links. "Area 0" is the name commonly given to the OSPF area with which all other OSPF areas in a given domain share routers. So the top label, T5, that ABR1 applies is an ATM VCI. Also, since ATM hardware does not lend itself to label-stack "popping," we will assume that ABR2's TDP message to TR0 specified a VCI value, T7, that ABR2 treats as though it were an explicit-NULL label on the top of the label stack. When TR0 swaps top VCI values, it therefore replaces T5 with T7, which ABR2 accordingly discards without using it as an index into its forwarding table. It instead uses T6 for that purpose. As was explained above in discussing the OSPF-message sequence, T6 identifies ABR2's forwarding-table entry for E-ASBR-destined packets.

The replacement-label stack in ABR2's forwarding-table entry for E-ASBR-destined packets does not include a label contributed by the AS-External LSA generated by E-ASBR and forwarded by TR2. Even if ABR2 were to use that AS-External LSA in generating its forwarding-table entry for E-ASBR, the result would not be a label contribution, because, unlike the MPLS Label field in the AS-External LSA that ABR1 received, that field in the AS-External LSA that ABR1 received carried an explicit-NULL value. But ABR2 did not use the received E-ASBR-designating AS-External LSA in computing the entry in its forwarding table for E-ASBR, because E-ASBR would also have sent Area 2 routers, including ABR2, an LSA of the Router (as opposed to AS-External) type in which the Link State ID field also specified E-ASBR. In accordance with normal OSPF forwarding-table computation rules, an LSA of this type is given precedence over an AS-External LSA for computing forwarding-table entries. Since that Router LSA therefore effectively masked the AS-External LSA, receipt of that AS-External LSA did not result in installation of a label in the corresponding forwarding entry for E-ASBR.

In contrast, the E-ASBR-specifying AS-External LSA that ABR1 received was not similarly masked. ABR1 would not have received an E-ASBR-specifying Router-type LSA, because ABR2 would have filtered that such an LSA from E-ASBR as part of its function as an area border router. So ABR1 would instead have received a summary-type LSA specifying a range that included E-ASBR, and normal OSPF rules for route computation would have given precedence to the more-specific AS-External LSA over the less-specific summary LSA.

Since ABR2's forwarding-table entry for E-ASBR therefore has no label contribution from the AS-External LSA, ABR2 replaces label T6 with only a single-label label stack. That stack's label is T8, the label whose use router TR2 requested for E-ASBR-destined packets. When it receives that packet, TR2 uses label T8 to find its forwarding information, which indicates that the replacement label stack should be the empty stack. That is, E-ASBR has sent TR2 a TDP message whose NULL value specified that TR2 should 10 remove but not replace the top labels on packets forwarded in accordance with the forwarding information for E-ASBR. This exposes to E-ASBR the label, T3, associated with E-ASBR's forwarding-table location that tells E-ASBR how to forward packets to the D-containing address range defined 15 by prefix P. E-ASBR accordingly forwards the packet to destination D.

A review of the foregoing operation reveals that it requires much less of the available virtual-circuit-indicator space than conventional schemes do. To appreciate this, 20 consider the link between ABR2593 and TR0. Although FIG. 6 depicts only one outgoing from link TR0, there are often very many. Moreover, although we have here discussed only a single exterior destination, the typical backbone can easily be called upon to carry traffic for thousands 25 of them. The resultant number of combinations of links and exterior destinations, in addition to the many interior destinations for which they may need to allocate virtual-circuit identifiers, may well exceed the capacity of the available virtual-circuit-identifier space.

30 In the forwarding scenario just described, though, TR0 was not required to allocate a virtual-circuit identifier to external destination D (or prefix P) in order to forward D-destined data packets. Indeed, it was not required even to allocate a virtual-circuit indicator to any autonomous-system border router. All that was required of it was that it 35 allocate a virtual-circuit identifier to the area border router involved. This can reduce the number of required virtual-circuit identifier by orders of magnitude.

Moreover, the approach described here is highly compatible 40 with existing routing methods. It builds upon the area partitioning for which many OSPF domains are already configured. It does not require implementation of any new protocol; it employs OSPF, which already is widely used for advertising topological information.

Also, it can be introduced into an OSPF domain gradually. Although not all of its benefits are afforded if not all relevant 45 routers implement it, routers that do can interoperate with those that do not, because the present invention's use of the External Route Tag field will be transparent to routers that have not yet been updated to take advantage of its benefits.

And it can be implemented in ways that differ in detail from the particular embodiment described above. For example, reflection will reveal that the use of the Filter at 50 ABR flag is not mandatory. It is included above only because it may be desirable in some circumstances to specify that filtering should not occur. For example, a system configured differently from the example above may have host routes to the ASBRs installed in all domain routers; e.g., area border routers would not summarize ASBR routes. In such a case, an ASBR could send an AS-External LSA whose Link-state ID field specifies an external destination, whose Forwarding Address field contains an address of that ASBR, and whose External Route Tag field contains the 55 label associated by that ASBR with its forwarding-table entry for that external destination. Routers receiving that LSA would then label packets intended for the external

destination with a label stack in which one label is the one that the ASBR thereby requested in the AS-External LSA, and another is a label requested by the next hop in the route to that ASBR. In such a situation, it would be necessary that the ASBR's choice of label and forwarding address be communicated to all routers, not filtered by area border routers, and the don't-filter value in the Filter at ABR field could signify this.

The present invention meets the need for virtual-circuit-indicator conservation in a way that is simple, robust, and easy to implement. It thus constitutes a significant advance in the art.

What is claimed is:

1. An OSPF domain configured as a plurality of OSPF areas and comprising a plurality of OSPF-speaking label-switching routers, including at least one area border router and one autonomous-system border router, that maintain forwarding information bases containing forwarding entries, identified by respective forwarding labels, that contain information that the label-switching routers use for forwarding 15 packets toward respective destinations, at least some of the label-switching routers responding to reception of an AS-External LSA by interpreting part of the AS-External LSA's External Route Tag field as indicating whether the External Route Tag field contains a forwarding label and, if so, interpreting part of the External Route Tag field as a 20 forwarding label and installing that forwarding label as a replacement label in any forwarding entry based on that LSA for the destination that the AS-External LSA's Link-State ID field specifies.

2. An OSPF domain as defined in claim 1 wherein each said area border router interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if the LSA External Route Tag field's first bit is a zero.

3. An OSPF domain as defined in claim 2 wherein each said area border router interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if a format part of the AS-External LSA's External Route Tag field contains a predetermined label-container-indicating value.

4. An OSPF domain as defined in claim 3 wherein the format part of the External Route Tag field is seven bits in length.

5. An OSPF domain as defined in claim 3 wherein the forwarding label occupies a twenty-bit-long part of the LSA External Route Tag field.

6. An OSPF domain as defined in claim 1 wherein, in response to receipt, from one said OSPF area, of each of at least some received OSPF AS-External LSAs whose Link-State ID fields specify autonomous-system border routers and whose External Route Tag fields it interprets as containing a forwarding label, at least one said area border router sends into another OSPF area a filtered OSPF AS-External LSA, whose Link-State ID field specifies the autonomous-system border router specified by the received OSPF AS-External LSA, whose Forwarding Address field contains an address of that area border router, and whose External Route Tag field contains the forwarding label associated with that border router with its forwarding-table entry that contains forwarding information for packets destined for the given autonomous-system border router.

7. An OSPF domain as defined in claim 6 wherein said at least one area border router that sends filtered OSPF AS-External LSAs interprets part of the External Route Tag fields that it interprets as containing forwarding labels as

also including a filter-at-ABR field and sends filtered OSPF AS-External LSAs only in response to received OSPF AS-External LSAs in which the filter-at-ABR field contains a predetermined do-filter value.

8. An OSPF domain as defined in claim 7 wherein each said area border router interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if the LSA External Route Tag field's first bit is a zero.

9. An OSPF domain configured to provide a plurality of OSPF areas, including at least first and second OSPF areas, the OSPF domain including:

- A) a first autonomous-system border router, located in the first OSPF area, that transmits requests for use of a first label on packets to be forwarded by it to a given external destination outside that OSPF domain;
- B) an area border router, located in the second OSPF area, that transmits requests for use of a second label on packets to be forwarded by it to the first autonomous-system border router;
- C) a transit router, located in the second OSPF area, that transmits requests for use of a third label on packets to be forwarded by it to the area border router; and
- D) a second border router that receives the requests from the transit, area border, and autonomous-system border routers and responds to at least some data packets from an external source outside the OSPF domain and destined for the given external destination by forwarding those data packets into the OSPF domain with a forwarding-label stack that includes the first, second, and third labels.

10. An OSPF domain as defined in claim 9 wherein the area border router sends its request as an OSPF AS-External LSA whose External Route Tag field includes the first label.

11. An OSPF domain as defined in claim 10 wherein the area border router sends its request as an OSPF AS-External LSA whose Link-State ID field specifies the first autonomous-system border router and whose Forwarding Address field contains an address of the area border router.

12. An OSPF domain as defined in claim 10 wherein the area border router sends its request as an OSPF AS-External LSA in which part of the AS-External LSA's External Route Tag field is interpreted by the second autonomous-system border router as indicating that the External Route Tag field contains a forwarding label.

13. An OSPF domain as defined in claim 9 wherein the area border router maintains a forwarding information base containing forwarding entries, identified by respective forwarding labels, that contain information that it uses for forwarding packets toward respective destinations, responds to reception of an AS-External LSA by interpreting part of the AS-External LSA's External Route Tag field as indicating whether the External Route Tag field contains a forwarding label, and, if so, interpreting part of the External Route Tag field as a forwarding label and installing that forwarding label as a replacement label in any forwarding entry based on that LSA for the destination that the AS-External LSA's Link-State ID field specifies.

14. An OSPF domain as defined in claim 10 wherein the area border router sends its request in response to receipt, from outside the second OSPF area, of an OSPF AS-External LSA whose Link-State ID field specifies the autonomous-system border router and whose External Route Tag field it interprets as containing a forwarding label.

15. An OSPF-speaking label-switching router comprising:

- A) a forwarding information base containing forwarding entries, identified by respective forwarding labels, that

15

contain information that the router uses for forwarding packets toward respective destinations;

B) an interface for receiving AS-External LSAs; and

C) circuitry for responding to reception of an AS-External LSA by interpreting part of the AS-External LSA's External Route Tag field as indicating whether the External Route Tag field contains a forwarding label and, if so, interpreting part of the External Route Tag field as a forwarding label and installing that forwarding label as a replacement label in any forwarding entry based on that LSA for the destination that the AS-External LSA's Link-State ID field specifies.

16. A router as defined in claim 15 that interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if the LSA External Route Tag field's first bit is a zero.

17. A router as defined in claim 16 that interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if a format part of the AS-External LSA's External Route Tag field contains a predetermined label-container-indicating value.

18. A router as defined in claim 17 wherein the format part of the External Route Tag field is seven bits in length.

19. A router as defined in claim 17 wherein the forwarding label occupies a twenty-bit-long part of the LSA External Route Tag field.

20. A router as defined in claim 19 wherein the format part of the External Route Tag field is seven bits in length.

21. A storage medium containing instructions readable by a processor to configure the processor to function as an OSPF-speaking label-switching router that includes:

- A) a forwarding information base containing forwarding entries, identified by respective forwarding labels, that contain information that the router uses for forwarding packets toward respective destinations;
- B) an interface for receiving AS-External LSAs; and
- C) circuitry for responding to reception of an AS-External LSA by interpreting part of the AS-External LSA's External Route Tag field as indicating whether the External Route Tag field contains a forwarding label and, if so, interpreting part of the External Route Tag field as a forwarding label and installing that forwarding label as a replacement label in any forwarding entry based on that LSA for the destination that the AS-External LSA's Link-State ID field specifies.

22. A storage medium as defined in claim 21 wherein the router interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if the LSA External Route Tag field's first bit is a zero.

23. A storage medium as defined in claim 22 wherein the router interprets part of the AS-External LSA's External Route Tag field as indicating that the External Route Tag field contains a forwarding label only if a format part of the AS-External LSA's External Route Tag field contains a predetermined label-container-indicating value.

24. A storage medium as defined in claim 23 wherein the format part of the External Route Tag field is seven bits in length.

25. A storage medium as defined in claim 23 wherein the forwarding label occupies a twenty-bit-long part of the LSA External Route Tag field.

16

26. A storage medium as defined in claim 25 wherein the format part of the External Route Tag field is seven bits in length.

27. For communicating in an OSPF domain configured to provide a plurality of OSPF areas, including at least first and second OSPF areas, a method comprising:

- A) from a first autonomous-system border router, located in the first OSPF area, transmitting requests for use of a first label on packets to be forwarded by the first autonomous-system border router to a given external destination outside the OSPF domain;
- B) from an area border router, located in the second OSPF area, transmitting requests for use of a second label on packets to be forwarded by the area border router to the first autonomous-system border router;
- C) from a transit router, located in the second OSPF area, transmitting requests for use of a third label on packets to be forwarded by it to the area border router; and
- D) from a second border router, receiving the requests from the transit, area border, and autonomous-system border routers and responding to at least some data packets from an external source outside the OSPF domain and destined for the given external destination by forwarding those data packets into the OSPF domain with a forwarding-label stack that includes the first, second, and third labels.

28. A method as defined in claim 27 wherein the request from the area border router is sent as an OSPF AS-External LSA whose External Route Tag field includes the first label.

29. A method as defined in claim 28 wherein the request from the area border router is sent as an OSPF AS-External LSA whose Link-State ID field specifies the first autonomous-system border router and whose Forwarding Address field contains an address of the area border router.

30. A method as defined in claim 28 wherein the request from the area border router is sent as an OSPF AS-External LSA in which part of the AS-External LSA's External Route Tag field is interpreted by the second autonomous-system border router as indicating that the External Route Tag field contains a forwarding label.

31. A method as defined in claim 27 further including:

- A) maintaining at the area border router a forwarding information base containing forwarding entries, identified by respective forwarding labels, that contain information that the area border router uses for forwarding packets toward respective destinations;
- B) responding to reception of an AS-External LSA by interpreting part of the AS-External LSA's External Route Tag field as indicating whether the External Route Tag field contains a forwarding label; and
- C) if so, interpreting part of the External Route Tag field as a forwarding label and installing that forwarding label as a replacement label in any forwarding entry based on that LSA for the destination that the AS-External LSA's Link-State ID field specifies.

32. A method as defined in claim 28 wherein the request from the area border router is sent in response to receipt, from outside the second OSPF area, of an OSPF AS-External LSA whose Link-State ID field specifies the autonomous-system border router and whose External Route Tag field the area border router interprets as containing a forwarding label.

* * * * *